

REMARKS

In the Office Action, the Examiner rejected the claims under 35 USC §103. The rejections are fully traversed below. The claims have been amended correct various typographical errors. Claim 18 has been cancelled. Claims 1, 3, 5-17, and 19-61 remain pending.

Reconsideration of the application is respectfully requested based on the following remarks.

REJECTION OF CLAIMS UNDER 35 USC §103(a)

In the Office Action, the Examiner has rejected claims 1, 12, 14, and 15 under 35 USC §103(a) as being unpatentable over Blumenau et al, U.S. Patent No. 6,260,120, ('Blumenau' hereinafter) in view of Testardi, U.S. Pub. No. 2003/0140210 A1 ('Testardi' hereinafter). The Examiner has also rejected claims 29-48, 50-53, and 59-61 under 35 USC §103(a) as being unpatentable over Blumenau and Testardi. This rejection is fully traversed below.

Various embodiments of the invention support the virtualization of storage in a storage area network. This is accomplished through the use of one or more network devices capable of being placed in a data path between the hosts and the storage devices. As a result, neither the storage devices nor the hosts require additional software or hardware to support storage virtualization. Moreover, multiple network devices may simultaneously manage the virtualization of heterogeneous storage devices.

The pending claims implement storage virtualization, as claimed, on a per-port basis. In other words, selected ports of one or more network devices may implement virtualization functionality in hardware and/or software. Any number of ports on a switch can manage

virtualization of its own traffic. This allows a network's virtualization capacity to scale with the number of ports.

Blumenau relates to storage mapping and partitioning among multiple host processors in the presence of login state changes and host controller replacement. See title. A storage controller is programmed to define a respective specification for each host processor of a respective subset of the data storage to which access by the host processor is restricted, and each specification is associated with a host identifier stored in the memory. When the storage controller receives a data access request from a host processor, it decodes a host identifier from the data access request, and searches the memory for a host identifier matching the host identifier decoded from the request. Upon finding a match, the respective specification of the respective subset for the host processor is accessed to determine whether or not storage specified by the storage access request is contained in the respective subset. If so, then storage access can continue, and otherwise, storage access is denied. Preferably, the host identifier decoded from the request is a temporary address assigned by the network, and also stored in the memory in association with each respective specification is a relatively permanent identifier for the host processor. See Abstract.

As shown in Fig. 22 of Blumenau and described in col. 26, lines 1-24, a cached storage subsystem 250 implementing virtual ports 268 includes two port adapters 260 and 261, each having two physical ports. The port adapters are programmed to provide respective virtual switches linking their physical ports to a set of virtual ports. As shown in Fig. 7, the cached storage subsystem 20 appears to include a single network device having a cache memory 32 accessible by both port adapters. This is further emphasized in col. 9, lines 25-29, stating "in a preferred form of construction, the cache memory 32 is composed of dynamic RAM memory cards mounted in a card-cage or main-frame, and the port adapters and storage adapters are programmed micro-processor cards that are also mounted in the card-cage or main-frame."

As set forth above, Blumenau neither discloses nor suggests implementing storage virtualization, as claimed, on a network device. While Blumenau does disclose the mapping of LUNs to logical volume numbers (see col. 22, lines 60-67), Blumenau fails to disclose the claimed inventions as implemented in a network device, wherein the network device is a switch, router, iSCSI gateway, or other network node configured to perform a switching function.

The Examiner admits that Blumenau “fails to teach wherein (b), (c), and (d) are performed by logic dedicated to and implemented by said port of the network device.” The Examiner seeks to cure the deficiencies of Blumenau with Testardi.

It is important to note that the claims, as amended, recite a dedicated processor that is dedicated to and implemented by a port of the network device. Since a processor is dedicated to a port, virtualization of storage may be scaled with the number of ports. Testardi fails to disclose or suggest the use of a dedicated processor that is dedicated to and implemented by a port of the network device. As such, the combination of the cited references would fail to achieve the desired result. In view of the above, Applicant respectfully requests that the Examiner withdraw the rejection of claims 1, 12, 14, 15, and 29-48, 50-53, and 59-61 under 35 USC 103.

In the Office Action, the Examiner has rejected claims 3, 5-11, 13, 16, and 17 under 35 USC §103(a) as being unpatentable over Blumenau and Testardi, and further in view of Lo et al, U.S. Pub. No. 2002/0103943 (‘Lo’ hereinafter). This rejection is fully traversed below.

The Examiner admits that Blumenau and Testardi fail to teach a network device, where the virtual storage unit comprises a VLUN or other virtual representation of storage on a storage area network. The Examiner seeks to cure the deficiencies of Blumenau and Testardi with Lo.

Lo discloses a distributed storage management platform architecture. See title. Lo does disclose the concept of storage virtualization. See paragraphs 0037, 0239. The behavior disclosed in Lo is storage network router-based, rather than being housed in the hosts, or in the storage arrays/subsystems. See paragraphs 0247, 0048-0049.

In no manner does Lo disclose or suggest performing mapping functionality within a system implementing virtualization of storage on a per-port basis. In fact, Lo discloses that functionality other than conversion between frame formats be performed by other entities. For instance, col. 19, lines 38-52 disclose that storage commands are directed to a storage traffic engine, which handles storage commands. Moreover, storage traffic is sent to the

storage traffic block. In other words, these commands and storage traffic are not handled by a port. As such, Lo teaches away from performing or obtaining a virtual-physical mapping in a system implementing virtualization on a per-port basis. Accordingly, Applicant respectfully requests that the Examiner withdraw the rejection of claims 3, 5-11, 13, 16, and 17 under 35 USC 103.

In the Office Action, the Examiner has rejected claim 49 under 35 USC §103(a) as being unpatentable over Blumenau and Testardi, and further in view of Lo. This rejection is fully traversed below.

Lo teaches a network device, where the type of traffic is iSCSI. (See paragraph 0128). However, as set forth above, Lo teaches away from obtaining or performing a virtual-physical mapping on a per-port basis. Accordingly, Applicant respectfully submits that claim 49 is patentable over the cited references.

In the Office Action, the Examiner has rejected claims 18 and 19 under 35 USC §103(a) as being unpatentable over Blumenau and Testardi, and further in view of Latif et al, U.S. Patent No. 6,400,730 ('Latif' hereinafter). This rejection is fully traversed below.

The Examiner admits that Blumenau and Testardi fail to teach wherein (b), (c), and (d) are performed by a processor dedicated to said port of the network device." The Examiner seeks to cure the deficiencies of Blumenau and Testardi with Latif. It is important to note that the limitation recited in claim 18 has been incorporated into the independent claims.

The Examiner cites col. 18, lines 8-42 of Latif, stating that Latif teaches a network device, wherein (b), (c), and (d) are performed by a processor dedicated to a port of the network device.

Col. 18, lines 8-42 of Latif indicates that the routing logic includes logic blocks that are dependent on the port type and other blocks that are common to different port types. Routing block is described as determining where a frame is routed based upon addressing

information within the data frame. Latif further indicates that address resolution logic can be shared by two different port interfaces. Moreover, Latif further indicates that additional logic can also be shared. For instance, Latif indicates that routing logic block 350 is shared by two port interfaces. Thus, Latif does not require that the routing logic be dedicated to a single port, but rather teaches that routing logic is to be shared among multiple ports. It is important to note that Latif specifically states that the “logic within Routing Logic block 350 can be implemented as hard coded logic or as a programmable method using a network processor.” Therefore, Latif implies that the routing logic that is shared by two port interfaces be implemented by a single processor. As such, Latif teaches away from each port having its own dedicated processor. It is also important to note that the cited portion of Latif does not explicitly require that the routing logic include claimed steps (b), (c), and (d). Accordingly, Applicant respectfully submits that the independent claims incorporating the limitation of claim 18 are patentable over the cited art.

With respect to claim 19, the Examiner further cites Latif, col. 17, line 34-col. 18, line 7. Claim 19 recites, in part, “at least one of the processor or the memory being further adapted for requesting a lock of the one or more physical storage locations by said port of the network device prior to submitting a read or write command to the one or more physical storage locations.” However, Applicant was unable to find a reference to the port requesting a lock of one or more physical storage locations prior to submitting a read or write command. In fact, the cited portion of Latif neither discloses or suggests requesting a lock in any manner. As such, the combination of the cited references would fail to achieve the desired result. Accordingly, Applicant respectfully submits that claim 19 is patentable over the cited art.

The Examiner has also rejected claims 21-28 under 35 USC §103(a) as being unpatentable over Blumenau, Testardi, Latif, and further in view of Brewer et al, U.S. Patent No. 6,876,656, (‘Brewer’ hereinafter). This rejection is fully traversed below.

As claimed in claims 20-28, a port may submit a lock request to another “master” port that manages all lock requests. This master port may also notify the requesting port when a lock request has been granted. The master port may similarly process “lock release” requests.

With respect to claim 21, Blumenau fails to disclose sending a lock request to a master port of a network device. Rather, Blumenau appears to disclose that locking and unlocking is performed by the cached storage subsystem, rather than contacting a master port of a network device.

The Examiner cites Brewer, asserting that Brewer teaches a method, wherein requesting a lock of the one or more physical storage locations comprises: sending a lock request to a master port of a network device within the storage area network, wherein the master port is adapted for managing lock requests. The Examiner cites the abstract and col. 8, lines 20-60 of Brewer. However, Applicant was unable to find any reference to requesting a lock or sending a lock request to a master port. Rather, it appears that these portions of Brewer disclose redirection rather than sending a lock request (e.g., to prevent other ports or network devices from accessing or altering the locked storage locations). Redirection is an entirely different concept from locking.

In addition, although Brewer does disclose sending a redirection command to a master port on the switch, the master port copies the redirection command to all other ports. In contrast, the master port of the pending claims manages all lock requests. In other words, a lock request is not copied to other ports. As a result, Brewer teaches away from the claimed invention. Moreover, the cited references would fail to achieve the desired result. In fact, the combination of the cited references would result in unnecessary copying of commands by a master port. As such, Brewer fails to cure the deficiencies of Blumenau, Testardi, and Latif. Accordingly, Applicant respectfully submits that claims 21-28 are patentable over the cited references.


Applicant respectfully submits that the independent claims are patentable over the cited references, separately or in combination. The dependent claims depend from one of the independent claims and are therefore patentable for at least the same reasons. However, the dependent claims recite additional limitations that further distinguish them from each of the cited references. The additional limitations recited in the independent claims or the dependent claims are not further discussed, as the above discussed limitations are clearly sufficient to distinguish the claimed invention from the cited reference. Thus, it is respectfully requested that the Examiner withdraw the rejection of the claims under 35 USC §103(a).

SUMMARY

If there are any issues remaining which the Examiner believes could be resolved through either a Supplemental Response or an Examiner's Amendment, the Examiner is respectfully requested to contact the undersigned attorney at the telephone number listed below.

Applicants hereby petition for an extension of time which may be required to maintain the pendency of this case, and any required fee for such extension or any further fee required in connection with the filing of this Amendment is to be charged to Deposit Account No. 50-0388 (Order No. ANDIP003).

Respectfully submitted,
BEYER WEAVER LLP



Elise R. Heilbrunn
Reg. No. 42,649

BEYER WEAVER LLP
P.O. Box 70250
Oakland, CA 94612-0250

Tel: (510) 663-1100